

Security without Identification

Card Computers to make Big Brother Obsolete

by [David Chaum](#)

You may soon use a personal "card computer" to handle all your payments and other transactions. It can protect your security and privacy in new ways, while benefitting organizations and society at large.

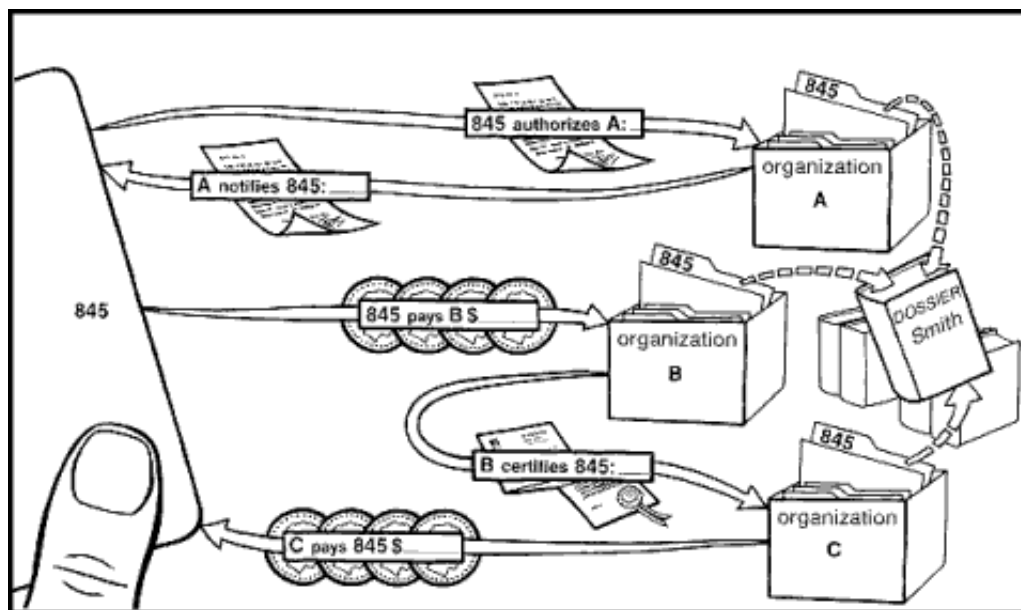
Computerization is robbing individuals of the ability to monitor and control the ways information about them is used. Already, public and private sector organizations acquire extensive personal information and exchange it amongst themselves. Individuals have no way of knowing if this information is inaccurate, outdated, or otherwise inappropriate, and may only find out when they are accused falsely or denied access to services. New and more serious dangers derive from computerized pattern recognition techniques: even a small group using these and tapping into data gathered in everyday consumer transactions could secretly conduct mass surveillance, inferring individuals' lifestyles, activities, and associations. The automation of payment and other consumer transactions is expanding these dangers to an unprecedented extent.

Organizations, on the other hand, are attracted to the efficiency and cost-cutting opportunities of such automation. Moreover, they too are vulnerable, as when cash, checks, consumer credit, insurance, or social services are abused by individuals. The obvious solution for organizations is to computerize in ways that use more pervasive and interlinked records, perhaps in combination with national identity cards or even fingerprints. But the resulting potential for misuse of data would have a chilling effect on individuals. Nevertheless, this is essentially the approach of the electronic payment and other automated systems now being tried. Although these systems will require massive investment and years to complete, their underlying architecture is already quietly being decided and their institutional momentum is growing.

This momentum is driving us toward a seemingly irreconcilable conflict, between organizations' need for security and the benefits of automation on one side, and individuals' need for ensured privacy and other protections on the other. But this conflict may be avoided by early adoption of a fundamentally different approach to automating transaction systems. This new approach is mutually advantageous: it actually increases organizations' benefits from automating, including improved security, while it frees individuals from the surveillance potential of data linking and other dangers of unchecked record keeping. Its more advanced techniques offer not only wider use at reduced cost, but also greater consumer convenience and protection. In the long run, it holds promise for enhancing economic freedom, the democratic process, and informational rights.

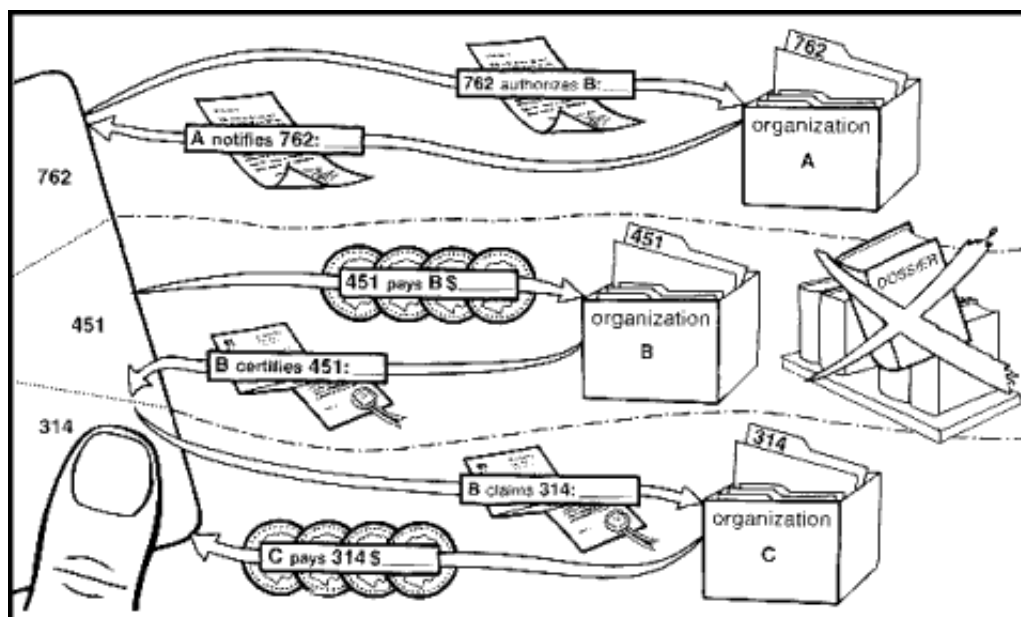
The New Approach and How it Differs

Three major differences define the new approach. The first is in the use of identifying information. Currently, many Western countries require citizens to carry documents bearing universal identification numbers. Drivers' licenses are being upgraded to perform a similar function in the United States, and efforts toward machine-readable national identity documents are expanding internationally. Meanwhile, organizations routinely use such essentially identifying data as name, birthday, and birthplace or name and address to match or link their records with those of other organizations.



UNIVERSALLY IDENTIFYING NUMBERS or other equivalent identifying information is presented by the individual card holder to each organization—in the current approach. Unrelated generic examples are shown of three kinds of transactions: *commutation*, in which the individual sends an authorizing message and receives a notifying message; *payment* in which the individual pays an organization or receives a payment; and *credential*, in which a certification that an individual has some credential is transferred from an organization B to an organization C. The identifying information—845—allows all transactions to be linked together into a dossier on the individual.

Under the new approach, an individual uses a different account number or "digital pseudonym" with each organization. No other identifying information is used. A casual purchase at a shop, for example, might be made under a one-time-use pseudonym; for a series of transactions comprising an ongoing relationship, like a bank account, a single pseudonym would be used repeatedly. Because of the input individuals have into the process by which the pseudonyms are created, they are ensured that their pseudonyms cannot be linked. This input also yields them the exclusive ability to use, and authenticate ownership of, their pseudonyms. Organizations too can protect themselves through their participation in forming the pseudonyms; among other safeguards, they can limit individuals to one pseudonym per organization and ensure that individuals are held accountable for abuses created under any of their pseudonyms.



DIFFERENT NUMBERS OR DIGITAL PSEUDONYMS are used with each organization by a personal card computer that the individual completely controls—under the new approach. The credential transfer is no longer just between organizations: it must now go through the card where the pseudonym, 451, used with the issuing organization 3 is transformed to the pseudonym, 314, used with the receiving organization C.

Systems using this approach can provide organizations with improved protection against abuses by individuals, and also allow individuals to ensure that pseudonyms cannot be traced across the dashed boundary lines, thereby preventing dossier compilation.

A second difference is in whose mechanism is used to conduct transactions. Today, individuals hold a variety of "tokens" issued to them by organizations. These range from traditional paper documents to plastic cards with magnetic or optical stripes or even embedded microcomputers. Such tokens are usually owned by the issuing organization and contain information that the individual holder can neither decipher nor modify. With the spread of automatic teller and point-of-sale terminals, individuals are being asked to perform more transactions directly using computer-controlled equipment. These terminals, and even the microcomputers in some current tokens, are physically tamper-resistant and contain secret numeric keys that securely code their communication with central computers. Individuals derive little direct benefit from these security provisions, however: in using such a transaction mechanism, they must take on faith the information it displays to them while revealing their own secrets to it.

With the new approach, an individual conducts transactions using a personal "card computer." This might resemble a credit-card-sized calculator and include a character display, a keyboard, and a short-range communication capability (like that of a television remote control). Such computers can be bought or even constructed, just like any other personal computer; they need have no secrets from, or structures unmodifiable by their owners. They can also be as simple to use as automatic teller machines. During a purchase at a shop, for example, equipment at the point of sale transmits a description of the goods and cost to the card, which displays this information to its owner. The card owner allows the transaction simply by entering a secret authorizing number on the card's keyboard. This same number is used by the owner to allow each transaction; without it, a lost or stolen card computer would be of very little use. A lost card's full capabilities, however, could be readily installed in a replacement, using backup data saved in a secure, encoded form at home or elsewhere.

The third defining difference is in the kind of security provided. Current systems emphasize

the one-sided security of organizations attempting to protect themselves from individuals, while the new approach allows all parties to protect their own interests. It relies both on individuals' card computers withholding secret keys from organizations and on organizations' computers devising other secret keys that are withheld from individuals. During transactions, the parties use these keys to form specially coded confirmations of transaction details, the exchange of which yields evidence sufficient to resolve errors and disputes.

The systems presented here for the new approach depend on currently used codes to secure organizations against abuses by individuals. Since these codes are "cryptographic," they can be broken, in principle, by trying enough guessed keys. Such guessing, however, is infeasible because of the enormous number of possible keys. In short, no proofs of security are known for these cryptographic codes. but nor are any feasible attacks. By contrast, the security card computers provide for individuals against the linking of their pseudonyms is "unconditional"-simple mathematical proofs can show that, with appropriate use of the systems, even collusion of all organizations and tapping of all communication lines could not Yield enough information to link the pseudonyms--regardless of how clever the attack or how much computation it uses.

In summary, if large scale automated systems or consumer transactions are actually to be built, the new approach offers a far more attractive way to structure them. Its specific advantages to individuals, organizations, and society at large will be argued further in the final section. The intervening three sections expand on its desirability and practicality for a comprehensive set of transaction types: communication, payments, and credentials.

Payment systems now being piloted for widespread use with the current approach include tamper-resistant card computers issued by banks and electronic connections between banks and retailers. The same basic mechanisms, however, could be designed to carry out payment transactions under the new approach. This in turn would allow new approach credential transactions to come naturally and gradually into use, with their applicability and benefits growing as computer and telecommunications infrastructures mature. The communication system proposed here would only begin to be practical with the advent of large-scale consumer electronic mail and would allow home use of the payment and credential systems. It is here presented first, however, since it most clearly illustrates some concepts central to the latter more immediately applicable systems.

Communication Transactions

As more messages travel in electromagnetic and digital form, it becomes easier to learn about individuals from their communication. Exposure of message content is one obvious danger, but this is already addressed by well-known coding techniques. A more subtle and difficult problem with current communication systems, however, is the exposure of "tracing information." An important kind of tracing information today is individuals' addresses, which organizations often require and which they commonly sell as mailing lists. The trend is toward greater use of such information. Comprehensive computerized data on who calls whom and when, for instance, are increasingly being collected and maintained by telephone companies. Electronic mail systems, some new telephone systems, and the proposed integrated services networks automatically deliver tracing information with each message. When such information is available on a mass basis, the pattern of each individual's relationships is laid bare. Furthermore, tracing information can be used as an identifier to link together all the records on an individual that are held by organizations with whom that

individual communicates. So long as communication systems allow system providers, organizations, or eavesdroppers to obtain tracing information, they are unsuitable for the new approach and, moreover, are a growing threat to individuals' ability to determine how information about themselves is used.

The other side of the issue is that current systems offer organizations and society at large inadequate protection against individuals who forge messages or falsely claim not to have sent or received messages. With paper communication, handwritten signatures are easily forged well enough to pass routine checking against signature samples, and they cannot be verified with certainty, even by expert witnesses. Also, paper receipts for delivery are too costly for most transactions, are often based solely on handwritten signatures, and usually do not indicate message content. As computerized systems come into wider use, moreover, the potential for abuse by individuals will increase. Solving these problems under the current approach might be attempted in several obvious ways: by providing recipients with the sender's address, by installing tamper-resistant identity-card readers or the like at every entry point to the communication system, and by keeping records of all messages to allow certification of delivery. But these security measures are all based on tracing information and thus are in fundamental conflict with individuals' ability to monitor and control information about themselves.

Both sets of problems are solved under the new approach. The nature of the solution is such that: individuals are able to send or receive messages without releasing any tracing information; receivers can show that messages were in fact sent to them, despite denial by the senders; senders can show that messages were in fact received, despite denial by the receivers; and message content is kept confidential. To make messages untraceable, a person's electronic mail computer conceals, in an unconditionally secure way, which messages it sends and receives. To prevent denial by a sender, each sender cryptographically codes messages in a way that each receiver can check, but that prevents anyone from being able to imitate the sender's coded "signature." These two concepts--untraceability and coded signatures--will recur intertwined in the payment and credential transaction types and are presented in separate subsections below.

Unconditional Untraceability

It is easy, in principle, to prevent a message sent by an organization from being traced to its individual recipient. The organization simply broadcasts all its messages to all individuals, and each individual's electronic mail computer then scans the broadcasts for messages addressed to any of its owner's pseudonyms. Thus only the individual's computer knows which of the broadcast messages its owner obtains.

Preventing a message sent to an organization from being traced back to its individual sender, however, requires some novel techniques; since any physical transmission can, in principle, be traced to its source. The concept of these techniques is illustrated by a hypothetical situation. Suppose two of your friends invite you to dine at a restaurant. After dinner, the waiter comes to your table and mentions that one of the three of you has already paid for the dinner--but he does not say which one. If you paid, your friends want to know (since they invited you), but if one of them paid, they do not want you to be able to learn which one of them it was.

The problem is solved at the table in the following simple way: Your friends flip a coin behind a menu so that they can see the outcome, but you cannot. It is agreed that each of them


will say the outcome aloud, but that if one of them paid, that one will say the opposite of the actual outcome. The uninteresting case is when they both say heads or both say tails: then everyone knows that you paid. If one of them says heads and the other says tails, however, then you know that one of them paid—but you have absolutely no information as to which one. You do know that the one you observed say tails paid if the coin toss was heads, and that the other one paid if the coin toss was tails. But since heads and tails tosses are equally likely, you learn nothing from your two friends' utterances about which one of them paid.

The system described allows the friend who paid to send you an unconditionally untraceable message; even though you know who says what, you cannot trace the "I paid" message, no matter how clever or prolonged your analysis.



UNCONDITIONALLY UNTRACEABLE MESSAGES are illustrated by a hypothetical situation (see text). The "I paid" message is unconditionally untraceable, since it can not be traced to a particular host—no matter how much computation or what approach is used.

This hypothetical system can be generalized and made practical (as detailed in reference [1]). One such generalization uses additional coins to allow more potential senders at the table, while preventing tracing even by collusion. Another breaks long messages into a sequence of parts, each of which is dealt with in a separate round of coin tosses and utterances. In practical communication systems, each participant's electronic mail computer would share secret numeric keys with other mail computers (just as hosts shared coin tosses behind their menus). Each mail computer then uses these keys to produce transformed sequences of digits (like a sequence of outcomes uttered at the table), which it sends through the mail network. The network combines all these transmissions to recover the original messages, which it broadcasts back to the mail computers (just as messages were audible and understandable to everyone at the table).



UNCONDITIONALLY UNTRACEABLE MESSAGES WITH NUMBERS are sent essentially as with words, except that everything is represented as 0's and 1's. Only the exclusive-or operation \oplus is used (defined as $1 \oplus 0 = 0 \oplus 1 = 1$ and $0 \oplus 0 = 1 \oplus 1 = 0$). The 0 or 1 outcome of the coin toss is shown as k . A host wishing to send the "I paid" messages which is represented as 1, transmits $k \oplus 1$; a host not wishing to send the message transmits only k . When the guest forms the exclusive-or of the two transmissions, [1] and [2], the result is 1. If one host sent the message and 0 if no host sent it — because k appears twice and cancels (since $k \oplus k = 0$ and $0 \oplus j = j$). If there are more hosts at the table, each flips a coin and shares the outcome with the host to the left, skipping the guest. Each host then forms a transmission as the exclusive-or of the two outcomes the host shares, exclusive-or'ed with an additional 1 if the host is sending the "I paid" message. Every coin toss appears twice and is cancelled in the exclusive-or that the guest forms from all the transmissions, and the result is again 1 if a host paid and 0 if no host paid. In actual computerized systems, real messages are encoded as sequences of 0's and 1's, and the whole protocol is repeated with new k 's for each digit to be sent. Senders noticing that their messages are being garbled by collision with other messages, wait randomly-chosen intervals before attempting to resend.

Digital Signatures

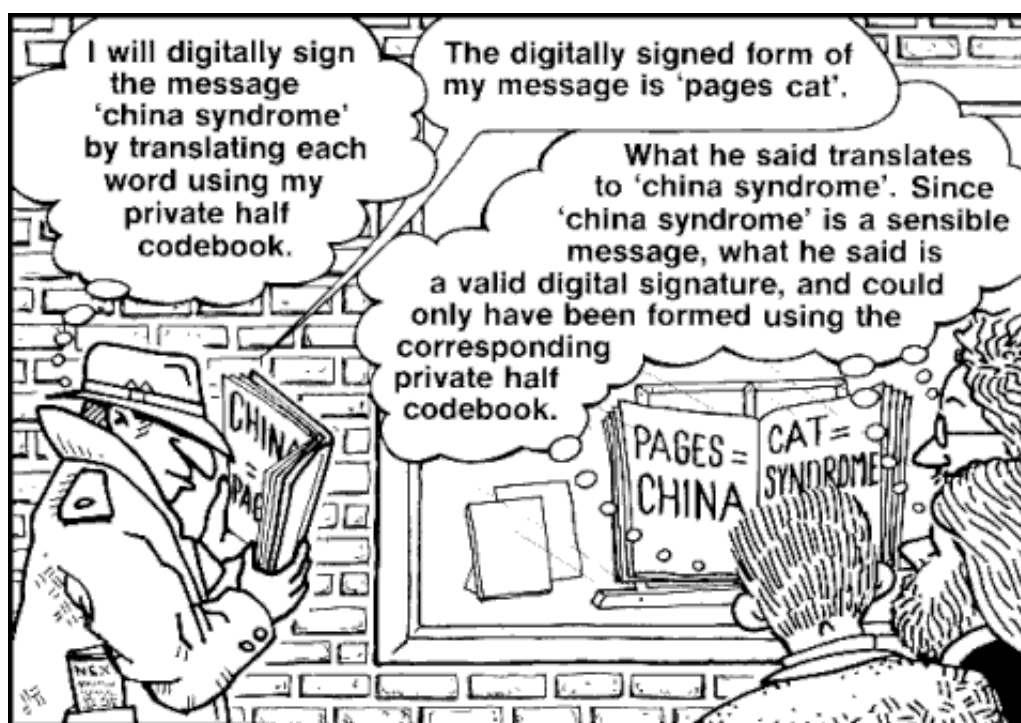
Now consider the problem of preventing senders from later disavowing messages they have sent. The solution is based on the concept of "digital signatures," which was first proposed by Diffie and Hellman [4]. To see how this concept works, imagine an old-fashioned codebook that is divided into two halves, like an English-French and French-English dictionary, except that only English words are used. Thus, if you look up an English word in the front half of the codebook, you find the corresponding (but usually semantically unrelated) English code word; if you then look this code word up in the back half, you find your original English word. Such codebooks are constructed by pairing off words at random: in the front half of the book, the pairs are ordered by their first words, and in the back half by their second words. For instance, if under "spy" the front half shows "why," then under "why" the back half shows "spy."

If you construct such a codebook, you can use it in your communication with an organization. You keep the front half as your private key, and you give the back half to the organization as your digital pseudonym with that organization. Before sending a message to the organization, you translate each word of the message into code using your private key; this encoded form

of the message is called a digital signature. When the organization receives the digital signature from you, it translates it back to the original English message using your digital pseudonym.

The immensely useful property of such digital signatures is their resistance to forgery. No one--not even the organization that has your digital pseudonym--can easily forge a digital signature of yours. Such forgery would entail creating something that your digital pseudonym decodes to a sensible English message. In the codebook analogy, of course, forgery merely requires searching through (or completely re-sorting by second words) the half of the book that is your digital pseudonym. With actual digital-signature cryptographic techniques currently in use, however, forgery is thought to require so much computation as to be infeasible even for the fastest computers working for millions of years. If an organization cannot forge a digital signature of yours, then it cannot successfully claim that you sent it a message that you in fact did not send. A third-party arbiter would decide in favor of an organization only if the organization could show a digital signature that yields the disputed message when translated with your digital pseudonym. But, because forgery is infeasible, the organization could obtain such a digital signature only if you had "signed" (i.e., encoded) the disputed message using your private key.

An organization could create its own private key and corresponding digital pseudonym (its own "codebook"); it would keep the private key (the front half) to itself, while widely disseminating the corresponding digital pseudonym (the back half). It would then use this private key to transform messages into digital signatures before sending them to individuals. The organization, unlike an individual, would create only a single private key and corresponding digital pseudonym, which it would use for all digital signatures it sends. Thus, anyone receiving a signed message from the organization would decode it using the organization's single, publicly disseminated digital pseudonym (commonly called a "public key"). These signatures would allow individuals to convince the organization, or anyone else if necessary, that the message had in fact been sent by the organization. In the payment and credential systems introduced in the following sections, such digital signatures formed by organizations play an important role.




DIGITALLY SIGNED MESSAGES are illustrated by a hypothetical use of old-fashioned codebooks (see text). Actual computerized digital signature systems now in use are not unconditionally secure, though the amount of computation required for forgery is thought to be unobtainable in practice.

Digital Signatures in Practice

Actual digital signatures are realized using numbers, and can be adapted to keep message content confidential and to certify delivery.

Practical, computerized digital-signature techniques work just as in the codebook analogy above, except that everything is done with two hundred-digit numbers. Each private key, and each digital pseudonym, is represented as one such number (rather than as a half codebook); each unsigned message and each signature is also represented as such a number (rather than as a string of English words). A standard, publicly available mathematical procedure lets anyone use a private key to form a corresponding digital signature from a message; a similar procedure allows anyone to recover the original message using the matching digital pseudonym (just as the simple procedure for looking up words in either half of the codebook can be public, so long as the private key is not). Another public mathematical procedure allows anyone to create a private key and corresponding digital pseudonym from a random starting point (just as the two halves of a codebook could be generated from a random pairing of words). Rivest, Shamir, and Adleman [5] proposed such a numeric digital-signature technique, which seems to be highly secure against forgery and could underlie the systems presented here.

Messages are kept confidential during transmission by using digital pseudonyms and private keys in a different way: before transmitting a message, the sender first signs it and then encodes the result using the digital pseudonym of the intended recipient. Thus, the signed message can be recovered only by decoding the transmission using the intended recipient's private key.



DIGITAL SIGNATURES WITH NUMBERS use special arithmetic systems, in which raising a number to a power \bar{x} scrambles it, and raising to a corresponding power \bar{x} unscrambles it: $(m^{\bar{x}})^x = m$. (The power \bar{x} acts

as the private half codebook, and the other power x acts as the corresponding half.) First the message is encoded as a onehundred-digit number, and then the digits are repeated to form a two-hundred-digit number \bar{m} with this special repeated-halves property. Next the signer raises the special number to a private power \bar{x} and makes the result known to others in transmission [1]. Someone receiving this digitally-signed message merely raises it to the corresponding digital-pseudonym power x and checks that the result has the special repeated-halves property. If it does, then the recipient knows that the message was signed by the holder of the corresponding private power

One way to protect against recipients falsely claiming not to have received messages is similar to the way paper mail is certified: messages are only given to recipients once they provide digitally signed "receipts" of delivery. Another method holds people responsible for messages that are made a matter of public record, like legal notices in newspapers. Since, under the new approach, messages are broadcast, they can be certified in this way at little additional expense. (A more fundamental advantage of making messages a matter of record is that it becomes easy to disprove false attributions of signatures--even if signatures could somehow be forged.) When this method is used with messages encoded for confidentiality, either party can display the signed message and point to the corresponding doubly encoded transmission in the public record as evidence that the message was available for receipt, since decoding the signed message with the digital pseudonym of the sender yields the message content, and encoding it with the pseudonym of the recipient yields the transmission in the public record.

Payment Transactions

The computerization of payments is giving payment system providers and others easy access to extensive and revealing information about individuals through payments made for purchases from shops, subscriptions, donations, travel, entertainment, professional services, and so on. Today, many paper records of when, how much, from whom, and to whom payment was made are translated into electronic form. The trend is toward capturing this payment data electronically, right at the point of sale. This facilitates the electronic capture of the potentially more revealing details of what was purchased. Moreover, computerization is extending the data capture potential of payment systems in other ways. One is through emerging informational services like pay television and videotex; another is through new systems that directly connect central billing computers to things like electric-utility meters and automobile identification sensors buried in toll roads. Just as, in communication systems, tracing information links all of an individual's records with organizations, payment data containing an account identifier links all of an individual's relationships involving payments.

From the other perspective, it is widely held that uncollectible payments made by consumers, such as credit card misuse and checks drawn against insufficient funds, cost society billions of dollars a year. Paper banknotes are vulnerable to counterfeiting and theft, and their lack of auditability makes them convenient for illicit payments such as bribes, extortion, and black-market purchases. Limiting all these abuses while automating seems to call for highly pervasive and interlinked systems that capture and retain account identifiers as well as other payment data--which is in clear conflict with the interests of individuals.

The nature of the new approach's solution to these problems ensures that organizations, even colluding with the payment system provider who maintains the accounts, cannot trace the flow of money between accounts. But the system provider does know the balance of each account, and if funds were to be transferred between accounts instantaneously, the simultaneous but opposite changes in balance would make tracing easy. Such tracing is


prevented because funds are withdrawn, held, and paid as multidenominational notes, in some ways like "unmarked bills." These notes are unlike paper banknotes, however, in that individuals, but not organizations, can allow transfers to be traced and audited whenever needed; this makes the notes unusable if stolen, and unattractive for many kinds of illicit payments. The fully computerized systems introduced here offer practical yet highly secure replacements for most current and proposed consumer payment systems (as detailed in [2]).

Blind Signatures for Untraceable Payments

The new-approach payment systems are based on an extension of digital signatures, called blind signatures. This concept is illustrated by an analogy to carbon-paper-lined envelopes. If you seal a slip of paper inside such an envelope and a signature mark is later made on the outside, then when you open the envelope, the slip will bear the signature mark's carbon image.

Consider how you might use such an envelope to make a payment. Suppose that a bank has a special signature mark that it guarantees to be worth one dollar, in the sense that the bank will pay one dollar for any piece of paper with that mark on it. You take a plain slip of paper sealed in a carbon-lined envelope to the bank and ask to withdraw one dollar from your account. In response, the bank deducts one dollar from your account, makes the signature mark on the outside of your envelope, and returns it to you. You verify that your sealed envelope has been returned with the proper signature mark on it. Later, when you remove the slip from the envelope, it bears the carbon image of the bank's signature mark. You can then buy something for one dollar from a shop, using the signed slip to make payment. The shop verifies the carbon image of the bank's signature on the slip before accepting it.

Now consider the position of the bank when the slip is received for deposit from the shop. The bank verifies the signature on the slip submitted for deposit, just as the shop did, and adds a dollar to the shop's account. Because the signature verified, the bank knows that the slip must have been in an envelope that it signed. But naturally the bank uses exactly the same signature mark to sign many such envelopes each day for all of its account holders, and since all slips were "blinded" by envelopes during signing, the bank cannot know which envelope the slip was in. Therefore it cannot learn from which account the funds were withdrawn. More generally, the bank cannot determine which withdrawal corresponds to which deposit--the payments are untraceable.



UNTRACEABLE PAYMENTS are illustrated by an analogy to envelopes and carbon paper. The individual (or, in the computerized analogue, the card) seals a blank slip of paper and a facing piece of carbon paper in an envelope, and supplies it to the bank. The bank deducts one dollar from the individual's account, applies a "worth one dollar" signature (stamp) to the outside of the envelope, and returns the unopened envelope to the individual. Upon receiving this, the individual verifies the bank's validating signature. Before making payment sometime later, the individual removes the envelope and carbon, leaving only the signed slip of paper. When the shop receives the slip, it verifies the carbon image of the validating signature on the slip, and supplies it to the bank for deposit. After also verifying the slip's validating signature, the bank honours the deposit, since it knows the slip must have been in an envelope it signed. The bank does not, however, know which of the many envelopes that it signed contained the slip, and thus the bank cannot trace the slip to the Individual's account. In actual computerized systems, unless the individual allows tracing, withdrawals on one side of the dashed boundary line and payments on the other side of it are unconditionally untraceable to each other—even if the bank and all other organizations collude.

In actual computerized systems, both slips and envelopes are replaced by numbers, the bank's signature mark becomes a digital blind signature, and payments are unconditionally untraceable (as described later in this section). The protocols for transacting withdrawals and payments would of course be carried out automatically by the card computer; its owner would merely have to allow each transaction by entering the secret authorizing number.

Extending the Envelope Analogy

Using note numbers provides protections similar to those offered by check numbers today. Since the bank is unable to see into the envelopes, nothing is revealed to the bank by a randomly chosen note number written on the slip before it is signed. (Alternatively, the slip's unique, random paperfiber pattern could represent the note number.) Stolen notes should not be accepted by the bank once the individual who withdrew the funds reports their note numbers. When given these numbers, the bank can also attest to the accounts to which funds have been deposited. Such traceability at the payer's initiative would discourage the use of these systems in bribery, extortion, black market purchases, and other illicit payments: recipients of such payments risk having their accounts traced if they deposit the notes, and being apprehended or just discovering that the notes are worthless if they try to spend them.

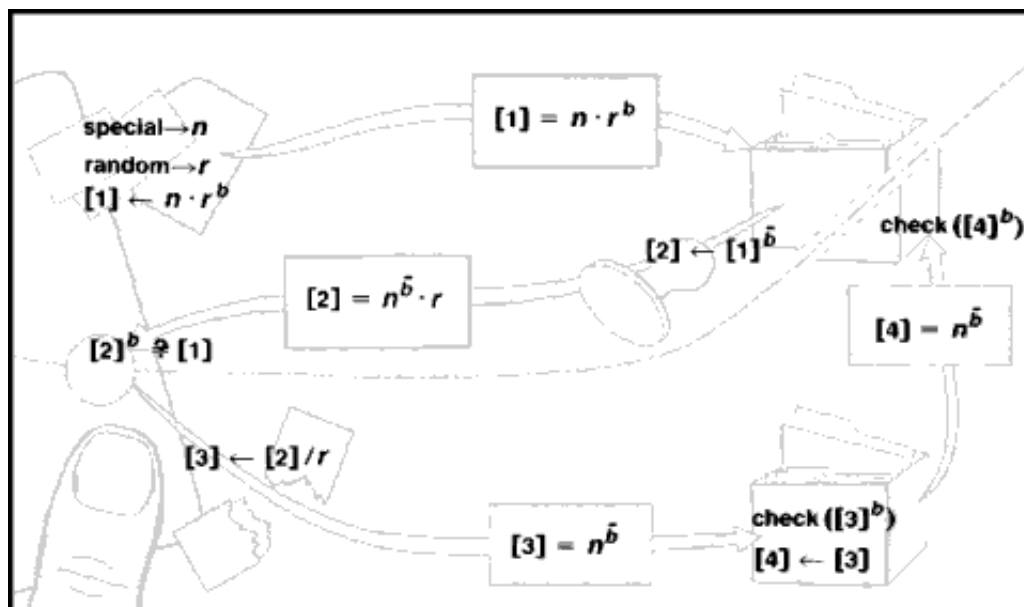
A variation prevents organizations (even colluding with banks) from tracing the accounts of individuals to whom they pay such things as wages, settlements, refunds, and rebates. The individual places a slip in an envelope as before and gives it to the paying organization, which

then supplies this blinded slip to the bank. The bank, without knowing which individual is involved, signs the envelope and charges the paying organization's account. Signed but still blinded, the slip is returned by the organization to the individual, who verifies the signature, and later removes the envelope and deposits the slip with the bank.

Other extensions to the basic concept offer replacements for today's payment systems attractive to both financial institutions and consumers. Regional clearing and signing centers would handle most of the work and responsibility for banks on a wholesale basis, while the banks could offer their own customized services. Different signatures would be used for different denominations. An adaptation allows routine transactions to be consummated in a way not requiring immediate or online interaction with a bank. Further variations permit the payment system to be used just as credit and debit cards are used today, with interest charges for credit and interest earnings on unspent debit-card balances.

Leaving the Analogy

Actual payment systems would work very much along the lines of the envelope analogy, except that they use no paper, only numbers. A note number is first created by a true random process within the individual's card computer (used like the random number or fiber pattern on the slip of paper). Next, the card computer transforms the note number into a numeric note that is the equivalent of the message: "This is note number: 59...2" (used like the slip of paper itself). The card computer then blinds this numeric note by combining it with a second random number (like the payer choosing an envelope at random and placing the slip in it). During withdrawal, the bank uses the private key of the desired denomination to form a digital signature on the blinded numeric note (like the signature mark made on the envelope). When the signed but still blinded note is returned, the card computer is able to unblind it by a process that removes the random blinding number from the digital signature while leaving the signature on the note (like the payer removing the envelope). Both the organization receiving payment and the bank use the bank's digital pseudonym to decode the signature; if the result is an appropriate message, this verifies the note's digital signature.



UNTRACEABLE PAYMENTS WITH NUMBERS are made much as in the paper analogy. First the individual's card computer chooses half the digits of n by a physical random process, and repeats these digits (actually in a scrambled form) to create the note number n with this special repeated-halves property (corresponding to choosing a suitable slip of paper at random in the analogy). The card also creates a

totally random number r (like choosing an envelope and carbon). The card then raises the random number r to the bank's "worth one dollar" public power b , multiplies this by the note number n (like sealing the slip in the envelope), and supplies the result to the bank in transmission [1]. The bank deducts from the account uses the corresponding private power \bar{b} to sign the transmission, and returns the result to the card in [2]. The card verifies that the bank returned exactly the right thing, and obtains the signed note by dividing out the random r (like removing the envelope and carbon). When a payment is made, the shop checks that transmission [3] is a signed special number, and then forwards a copy [4] to the bank for deposit. The bank checks the signature just as the shop did, and accepts the deposit if the valid note has not already been deposited. If Individuals do not divulge the random r 's their cards create, then the [1]'s are unconditionally untraceable to the [4]'s, since there is exactly one r that would make any [2] correspond with any [4].

A conceivable danger for the bank is that the same numeric note might be deposited more than once. To prevent this, a list of note numbers accepted for deposit is maintained and only note numbers not already on the list are accepted and recorded. The cost of maintaining such a list can be far less per transaction than the transaction cost of current payment systems, since expiration dates built into note numbers allow old numbers to be deleted from the list.

Another conceivable danger is that the bank's digital signature could be forged, which would allow counterfeiting. The security against this kind of threat is based on the underlying digital-signature cryptographic technique, which is currently being proposed as an international standard and is already used by banks and even by nuclear agencies. The odds of someone guessing a valid, signed numeric note, or of any two independently chosen note numbers being the same in the foreseeable future, are less than 1 in 10 to the 75th power.

The numeric notes are unconditionally untraceable: the bank cannot learn anything from the numbers about the correspondence between withdrawals and deposits. In the hypothetical restaurant situation, both outcomes of each coin toss were equally likely, which meant that every correspondence between senders and messages was equally likely. Similarly, because all suitable numbers are equally likely to be used for the independent blinding of each note, all correspondences between withdrawals and deposits are equally likely.

Credential Transactions

In their relationships with many organizations, there are legitimate needs for individuals to show credentials. The term "credentials" is used here to mean statements concerning an individual that are issued by organizations, and are in general shown to other organizations. In the past, credentials primarily took the form of certificates like passports, driver's licenses, and membership cards. Before computerization, such certificates provided individuals with substantial control over access to their credentials, though the certificates also often revealed unnecessary and identifying information like address, birthdate, and various numbers. Today, such identifying information is being used to link records on certificate holders; it even allows them to be "blacklisted" or denied services because of reports from organizations that may be erroneous, obsolete, or otherwise inappropriate for the decision at hand. Where no substantiating certificate is required to be shown, as with application or tax forms, much similarly unnecessary or overly detailed information is demanded, presumably to allow confirmation. But confirmation itself can link further information and lead back to inappropriate records. The control over credential information that certificates once provided to individuals is thus being circumvented and rendered illusory by computerization.

The countervailing problem is that credentials are subject to widespread abuse by individuals, who can easily modify or copy many kinds of paper and plastic certificates with today's

technology. This is one reason why certificates are in effect being reduced to the role of providing identifying information, and organizations are maintaining the credentials themselves. To check on unsubstantiated credential information, organizations are also rapidly deploying so-called matching techniques, whereby they use identifying information to link and share records on individuals. Many organizations may also need the ability to blacklist individuals or to determine whether they are already blacklisted. As the number of such organizations grows, certificates or even matching techniques become impractical, hence the creation of large centralized databases on individuals. The use of multiple complete identities by sophisticated criminals is a related problem. As with communication and payments, the obvious countermeasures under the current approach--widespread use of highly secure identity documents linked to centrally maintained credentials--are in direct conflict with individuals' ability to determine how information about themselves is used.


With the new approach's solution, an individual can transform a specially coded credential issued under one pseudonym into a similarly coded form of the same credential, which can be shown under the individual's other pseudonyms. Since these coded credentials are maintained and shown only by individuals, they return control similar to that formerly provided by certificates; and since they are convenient to use, they obviate the need for unsubstantiated credentials and for matching. Individuals can also tailor the coded form they show to ensure that only appropriate information is revealed or used to make particular decisions, and can ensure that obsolete information becomes unlinkable to current pseudonyms. Abuses of credentials by individuals, such as forgery and improper modification or sharing, are prevented by the cryptographic coding and the protocols for its use. Since each person is able to have at most one pseudonym with any organization requiring such protection, multiple complete identities are also prevented. Moreover, accountability for abuses perpetrated under any of an individual's pseudonyms can still be assured, without the need for centralized databases.

The Basic Credential System

The essential concept is again illustrated by analogy to carbon-lined envelopes, only this time the envelopes have windows. First, you make up numeric pseudonyms at random and write them on a plain slip of paper. When you want to get a credential from an organization, you put the slip in a carbon-lined envelope with a window exposing only the pseudonym you use with that organization. Upon getting the envelope from you, the organization makes a special signature mark in a repeating pattern across the outside of it, and the carbon lining transfers the pattern to the slip. This signature pattern is the credential; the type of pattern corresponds to the kind of credential the issuing organization decides to give you, according to the pseudonym they see through the window. When you get the envelope back from the issuing organization, you verify the credential signature pattern. Before showing the credential to another organization, you place the slip in a different envelope with a window position that exposes only the pseudonym you use with that organization, along with some of the adjacent credential signature pattern. The receiving organization can verify, through the window, the pseudonym you use with it as well as the signature pattern. In this way, you can obtain and show a variety of credentials.


An organization can ensure that no individual is able to transact with it under more than one pseudonym. One way an individual could attempt to use more than a single pseudonym with an organization is to use different pseudonyms on the same slip of paper. This is prevented by a standard division of the slip into positional zones, such that each zone is assigned to a

particular organization; an envelope is accepted by an organization only if the window position exposes that organization's zone, bearing a single indelibly written pseudonym. A second way of attempting to use more than one pseudonym per organization is to use more than one slip. This is prevented by the establishment of an agency that issues a single "is-a-person" credential signature to each individual. Other organizations accept only envelopes with this signature recognizable through the window. The agency ensures that it issues no more than one signature per person by taking, say, a thumbprint and checking that the print is not already on file before giving the signature. This collection of prints poses little danger to individuals, however, since the prints cannot be linked to anything.



UNTRACEABLE CREDENTIAL TRANSFERS BETWEEN PSEUDONYMS are illustrated by an analogy to window envelopes and carbon paper. The individual writes the pseudonyms on a slip and seals it, along with a facing piece of carbon paper, in an envelope the window of which exposes only the pseudonym 523 used with organization X. Then X applies a signature (stamp) on the outside of the envelope received, having chosen C as the repeating pattern that indicates the kind of credential Issued. The individual verifies the signature returned. When the individual later wishes to show the credential to organization Y, the original envelope and carbon are discarded, and the slip is placed in a new envelope the window of which exposes only the pseudonym 965 used by the individual with Y. Now Y verifies the signature through the window of the envelope and knows that 965 has been issued credential C. Organization Y cannot however, learn the other pseudonyms written on the slip. Actual computerized systems maintain the unconditional untraceability of pseudonyms across the dashed boundary line.

The pseudonyms used by individuals are untraceable, in the sense that envelopes give no clue, apart from the signatures shown, about the other randomly chosen pseudonyms they contain. Actual systems based on card computers would provide unconditional untraceability using digital blind signatures on numbers (as detailed in [3]).



UNTRACEABLE CREDENTIALS WITH NUMBERS also follow the paper analogy. The so-called "one-way" function f is easily computed by a publicly known procedures, but its inverse is thought to be infeasible to compute. Organization X determines the validity of both transmissions it receives, [1.1] and [1.2], by verifying that the first is a signature on the one-way function of the second. Later, X provides the signature for the desired credential on [1.2]. The card verifies the signature received as [2], and replaces s_x by s_y . Organization Y verifies [3.1] and [3.2] just as X did for [1.1] and [1.2]. Upon receiving the credential [3.3], Y verifies that it is a signed copy of [3.2].

Not illustrated is how a special organization Z ensures that the [1.1]'s (and [3.1]'s) are of the proper form, but does not obtain information useful in tracing. First the card supplies many candidates to Z, each of the form $q_n = f(u \cdot s_n^c) \cdot t_n^z$, where u is the special pseudonym used by the individual with Z, and $s_n = f(s_n')$ and $t_n = f(t_n')$, with s_n' and t_n' created at random by the card. When the card later learns which candidates q_n have been selected at random for inspection by Z the card supplies the corresponding s_n' and t_n' for each. This allows Z to verify that $q_n = f(u \cdot f(s_n')^c) \cdot f(t_n')^z$. If all inspected candidates verify, then Z supplies the signed form of all uninspected candidates. The card transforms a signed candidate into [1.1], for example, by dividing out t_n . (in one extension, the chance that improper candidates are used successfully can be kept negligibly small by Z returning signatures only on products of candidates. In another, using s 's as exponents with bases that are known generators having public signatures allows an unlimited number of signature types c_i .)

Revealing Only Necessary Information

You need not show all your credentials to every organization; you can restrict what you show to only what is necessary. Because of the way the credential signature patterns repeat across the slips, a recognizable part of each signature pattern appears adjacent to each pseudonym. To prevent certain credentials from being seen, though, you could simply black out parts of an envelope's window when showing it to an organization. But more flexible restrictions are possible using your card computer. It serves as the single database of all your credentials--and you alone control which queries from organizations it answers.

A typical such query might be: "Does the owner of pseudonym 72...4 have credentials sufficient to meet the requirement:...?" Your card can issue a convincing affirmative response only when it does in fact have credential signatures satisfying the requirement. But the card ensures--unconditionally--that organizations cannot learn any more about your credentials from its responses than the affirmations themselves. You might use it to convince an organization that your age, income, and education, for instance, meet their entry requirements

in at least one way, without revealing any more than just that fact. Or, when a survey requires credentials for substantiating responses, using a different pseudonym for each response ensures that no more is revealed than the total number of each type of response.

Actual queries and responses can be realized as follows: an organization encodes a new credential into the query message itself, in such a way that the credential can be decoded using any one of several qualifying combinations of other credentials as the key. If any qualifying combination is held, then this new credential can be decoded and shown to the organization as the response. It can also be retained for later use, which additionally permits the gradual replacement of older and more detailed credentials by more appropriate summary ones. When such query messages are made public so that everyone can use them, they provide for public and verifiable rules for decisions about individuals.

Some Uses of Credentials

The new approach supports most varieties of credentials used today. Some of these, like educational degrees, are lifelong, while others, like student cards, are valid only for prescribed periods. Still others, like membership cards, usually have long-term validity, but their certificates typically expire at the end of each year, thereby allowing their issuers to effectively revoke the credential by withholding new certificates.

A less common but still used kind of credential allows organizations in effect to blacklist individuals, without maintaining a central list of identities. Suppose, for example, that credentials are issued for filing tax forms, so that each adult citizen should get such a credential every year. Organizations might routinely modify their queries to include the requirement that adult citizens have filed tax forms for the last year. This would blacklist those who had not complied by barring them from relationships with organizations.

In actual widespread use, where many organizations may occasionally need to blacklist some individuals, such a mechanism is neither practical nor desirable: queries would have to demand vast numbers of credentials, while individuals would be unable to protect themselves against being blacklisted by organizations even with which they have had no contact.

Authorized Blacklisting Without Lists

These problems of wider use can be solved by techniques that require an organization to obtain, directly from an individual, the authorization to blacklist that individual for a specified reason. Organizations would insist on such authorizations as are appropriate before establishing or extending relationships.

The way these techniques work is illustrated by applying the envelope analogy to buying goods on credit. A special row of zones is reserved on each slip for this purpose. You provide the shop where you make the credit purchase with an envelope that has (in addition to any window you may ordinarily use with that shop) a window exposing one of these reserved zones. The shop first broadcasts the numeric pseudonym it sees indelibly written in that reserved zone, so that when no other organization objects, the shop is assured exclusive use of that zone.

When you later pay the shop, it gives you a resolution credential signature mark; unlike the credential signature marks previously described, it is made only on the single zone to which it

applies. If some of the reserved zones remain unused, you can show them to a "voiding" agency that obtains exclusive use of these unneeded zones in the same way as do shops, and then issues a resolution signature mark on each.

Only when you repay by deadline all due loans can you obtain resolution signature marks on each zone of the reserved row. Then you can demonstrate that you are not blacklisted, without revealing more, just by showing that all of your reserved zones have their resolution signatures. You do this by presenting an envelope that has a slit-shaped window positioned over the reserved row. It exposes only a narrow band of each reserved zone's resolution credential signature, while concealing the pseudonym-bearing parts of the zones that were shown separately to lenders or the voiding agency. In actual systems, card computers would obtain and show digital signatures for this purpose as part of their general management of the reserved row.

Preventing Use of Untimely Information

The mechanisms of the new approach can both guarantee individuals time to review credential information before it is required, and unconditionally ensure them the ability to shed such information once it is outdated.

If individuals can expect to receive their resolution credentials some "cooling-off" interval before they are needed, instead of at the last minute, then there may be time to resolve errors or disputes before any unnecessary consequences occur. Organizations may not wish to increase the maximum delay before blacklisting takes effect, but some cooling-off interval can always be provided without doing so. For example, when a different resolution credential is valid for each calendar month and organizations provide them just before the beginning of the month, then the maximum delay before blacklisting takes effect is one month and there is no cooling-off interval. But this same maximum delay can be maintained while providing cooling-off intervals half a month long: twice a month, organizations issue credentials that expire a month after their issue date, so that a credential remains valid for a half-month interval following the scheduled issue of its successor.

If individuals change pseudonyms periodically, they cannot be linked to obsolete information. The initial information associated with new pseudonyms would be provided through the transfer of credentials from previous pseudonyms. The changeovers could be staggered to allow time for completion of pending business.

There are additional benefits to changing pseudonyms beyond the weeding-out of obsolete information. For one thing, the periodic reduction to essentials prevents organizations from gradually accumulating information that might ultimately be used to link pseudonyms. Moreover, for individuals to be able to transfer all the initial information for a period, they must know each organization's information demands, they must know where each piece of information comes from, and they must consent to each such transfer. Information linkable by each organization is thus known to and agreed on by individuals--that is, individuals can monitor and control it.

Micro- and Macrocomparisons

Advantages to Individuals

As the public becomes more aware of the extent and possibilities of emerging information technology, there should be a growing demand for the kinds of systems described here. Compared to the current approach, individuals stand to gain increased convenience and reliability; improved protection against abuses by other individuals and by organizations; monitorability and control; and full access to transaction systems.

Increased convenience derives from the freedom of individuals to obtain their card computers from any source, to use whatever hardware or software they choose, and to interface with communication systems wherever they please. This permits card computers to be adapted to the requirements of sophisticated, naive, and handicapped users alike. The systems need be no more complicated to use than under the current approach. People might choose never to actually see their pseudonyms or to be concerned with other implementation details.

The individual is ensured reliable system access by a numeric key with which the card computer encodes backup copies of its contents, and which allows a replacement card to recover these contents. Since this key should be 40 or more digits in size, it might be impractical for its owner to remember. Known techniques allow the key to be divided into parts, each of which can be given to a different trustee. This provides certain subsets of the trustees with the ability to recover the key, while insufficient subsets would be unable to learn anything about it. Still other subsets, given parts of the owner's secret authorizing number, would be able to take over the owner's affairs when needed. These provisions are an example of how an individual's power to designate proxies, a power now enjoyed by organizations, is ensured.

Abuse of a lost or stolen card computer by another individual would be very difficult without the owner's secret authorizing number as asserted earlier. This is because the card would require the authorizing number, which might typically be about six digits long, before allowing transactions. A reasonably tamper-resistant device within the card computer could: read fingerprints or the like to prevent use by anyone but the card owner; accept a special authorizing number that the owner could use in case of duress to trigger a prearranged protective strategy; and permit only the current owner to reset the card for a new owner, to prevent its use as a replacement by a thief. Even if sophisticated criminals were to extract the card's information content and the owner were not to cancel in time using backup data, a great many guesses at the authorizing number might have to be tried with organizations before the actual number could be determined. This would make such attacks very likely to be detected and to fail.

The new approach protects individuals unconditionally from abuses by organizations, such as the false attribution of messages, and from organizations blacklisting without advance warning. Moreover, individuals are provided with secure relationships without ever having to sacrifice the protection of their pseudonyms by revealing linking information--but they can always do so if they choose. While it is relatively easy for individuals to provide convincing evidence only of their role in particular transactions, it is even possible for them to provide evidence that they were not involved in certain other transactions. For example, in communication transactions, individuals could show that their physical entry to the system was not used to send a particular message; in payment transactions, they could show that a payment did not involve their account; and in credential transactions, they could show that a pseudonym was not among the set obtainable under their thumbprint.

The primary way that individuals gain monitorability and control is through their ability to prevent linking. Some linking of separate relationships might occur if, for instance, a

consumer actually wanted to be recognized, or as part of an investigation or other exceptional situation. But the linking of some relationships does not, in general, allow others to be linked, and the regular changing of pseudonyms allows linkings to be shed over time. In addition, the scope of an individual's separate, unlinkable relationships need not depend on the legal or administrative structure of the organizations involved; an individual might use the same pseudonym with different organizations or, when allowed, different pseudonyms with the same organization. Naturally, the scope of relationships, along with such things as the level of detail in credentials and the frequency of pseudonym changeover, must be adjusted to provide the desired degree of protection against inference by statistical or pattern recognition techniques. Such protections would likely create a widespread expectation of control over information; thus, as similar expectations have done in the past, it might also engender commensurate legal safeguards.

Individuals would have the same access to systems as organizations, in addition to enjoying the same protections; such parity is precluded under the current approach in efforts to protect the security of organizations. A new-approach payment, for example, could be made between two friends using their card computers. A small business would even be able to handle all customer transactions, using only a card computer.

Advantages to Organizations

Organizations have an interest in cultivating the goodwill of individuals. But they gain further direct benefits from the advantages to individuals described earlier, since in making their own transactions, they have many of the same concerns as individuals. Moreover, the new approach offers them reductions in cost; reductions in the quantity and sensitivity of necessary data; and improved security against detectable, undetectable, and extrasystemic abuses.

The systems described here would be less costly for organizations than comparable systems based on the logical extension of the current approach. This is primarily because the latter requires widely trusted, tamper-resistant devices at all points of entry to transaction systems. Such a requirement implies substantial initial agreement, outlay, and commitment to design, and can be expected to result in technology that is outdated when systems come into widespread use. Furthermore, the tamper resistance techniques currently contemplated require significant compromise in security, even at high cost. The new-approach system provider need not supply user organizations with tamper-resistant terminal equipment for each entry point, any more than it need supply card computers to individuals. Thus, user organizations can supply their own terminal equipment wherever they please and take advantage of the latest technology. Although these cards and terminals make more sophisticated use of cryptographic techniques than does equipment envisioned under the current approach, this difference between the two is just a fraction of a chip in the technologies of the near future.

The new approach reduces the sensitivity and the quantity of consumer data in the hands of organizations; by the same token, it reduces their exposure to incidents that might incur legal liability or hurt their public images. Reductions in data could also streamline operations, and the increased appropriateness of the remaining data could provide a better basis for decision making. As electronic mail replaces paper mail, individuals' computers may routinely reject unsolicited commercial messages and instead seek out only desired information. Thus, data for targeting such messages might become superfluous even under the current approach. The

new approach's protections, however, may compensate by making individuals less reluctant to provide information for surveys and the like.

Under either approach, if an automated transaction system detects sufficiently serious abuse or default by an individual, the best it can do is to lock that individual out. This is because the individual can always step outside such a system's controls by "going underground." The new-approach systems can lock individuals out, but can also have a cooling-off interval built in to allow matters to be resolved before lockout is needed. The approach also reduces the need for such measures, however, since its mechanisms allow organizations or society at large the flexibility to set policy that establishes a desired balance between prior restraint, as in the basic payment system, and accountability after the fact, as with credit or other authorized blacklisting functions.

Undetectable abuse by individuals acting alone seems to be precluded by the systems of the new approach. But no transaction system is able to detect an individual who obtains something through legitimate use of the system and then transfers it to another person by some means outside the system. Transferring the ability to use a communication system to others is an instance of the proxy power already mentioned, which could be inhibited under the current approach. In the context of the payment system, such transfers can be treated as illicit payments, which are deterred by the use of note numbers. The credential system directly prevents the transfer of credentials from the pseudonyms of one person to those of another. Currently, "in-person" proxy is prevented by certificates bearing photos. Such photo tokens could still be used with the new approach, if and when needed; but they might include only a photo, an indication of the kind of credential, and possibly a digital pseudonym.

Meanwhile, it is too easy to step outside current transaction systems by using coin phones, sending anonymous letters, dealing in cash, and using false credentials. Significantly improved security, particularly against more sophisticated abuse, can only be obtained with comprehensive automated systems. But such systems under the current approach may meet with broad-based resistance from individuals--especially once they become aware of the alternatives posed by the new approach.

Implications for the Future

Large-scale automated transaction systems are imminent. As the initial choice for their architecture gathers economic and social momentum, it becomes increasingly difficult to reverse. Whichever approach prevails, it will likely have a profound and enduring impact on economic freedom, democracy, and our informational rights.

Restrictions on economic freedom may be furthered under the current approach. Markets are often manipulable by parties with special access to information about other participants' transactions. Information service providers and other major interests, for example, could retain control over various information and media distribution channels while synergistically consolidating their position with sophisticated marketing techniques that rely on gathering far-reaching information about consumers. Computerization has already allowed these and other organizations to grow to unprecedented size and influence; if continued along current lines, such domination might be increased. But the computerization of information gathering and dissemination need not lead to centralization: integrating the payment system presented here with communication systems can give individuals and small organizations equal and unrestricted access to information distribution channels. Moreover, when information about the transactions of individuals and organizations is partitioned into separate, unlinkable

relationships, the trend toward large-scale gathering of such information, with its potential for manipulation and domination of markets, can be reversed.

Attempts to computerize under the current approach threaten democracy as well. They are, as mentioned, likely to engender widespread opposition; the resulting stalemate would yield security mechanisms incapable of providing adequate prior restraint, thus requiring heavy surveillance, based on record linking, for security. This surveillance might significantly chill individual participation and expression in group and public life. The inadequate security and the accumulation of personally identifiable records, moreover, pose national vulnerabilities. Additionally, the same sophisticated data acquisition and analysis techniques used in marketing are being applied to manipulating public opinion and elections as well. The opportunity exists, however, not only to reverse all these trends, by providing acceptable security without increased surveillance, but also to strengthen democracy. Voting, polling, and surveys, for example, could be conveniently conducted via the new systems; respondents could show relevant credentials pseudonymously, and centralized coordination would not be needed.

The new approach provides a practical basis for two new informational human rights that is unobtainable under the current approach. One is the right of individuals to parity with organizations in transaction system use. This is established in practice by individuals' parity in protecting themselves against abuses, resolving disputes, conferring proxy, and offering services. The other is the right of individuals to disclose only the minimum information necessary: in accessing information sources and distribution channels, in transactions with organizations, and--more fundamentally--in all the interactions that comprise an individual's informational life.

Advances in information technology have always been accompanied by major changes in society: the transition from tribal to larger hierarchical forms, for example, was accompanied by written language, and printing technology helped to foster the emergence of large-scale democracies. Coupling computers with telecommunications creates what has been called the ultimate medium--it is certainly a big step up from paper. One might then ask: To what forms of society could this new technology lead? The two approaches appear to hold quite different answers.

Acknowledgements. The author is pleased to thank Jan-Hendrik Evertse, Wiebren de Jonge, and Ronald L. Rivest for discussions during the early development of some of the ideas herein presented, as well as everyone who showed interest in and commented on this work.

References

1. Chaum, D. The dining cryptographers problem: Unconditional sender and recipient untraceability. Available from the author.
2. Chaum, D. Privacy protected payments: Unconditional payer and/or payee untraceability. Available from the author.
3. Chaum, D. Showing credentials without identification: Transferring signatures between unconditionally unlinkable pseudonyms. Available from the author.
4. Diffie, W., and Hellman, M.E. New directions in cryptography. IEEE Trans. Inf: Theory, IT-22, (November 1976), 644-654.

5. Rivest, R., Shamir, A. and Adleman, L. A method for obtaining digital signatures and public-key cryptosystems. Communications of the ACM, 21, 2, (February 1978), 120-126.

Copyright 1987 by David Chaum. (A related earlier work appeared in Communications of the ACM, 28, 10, Oct. 1985; ACM is acknowledged.) Copies not made or distributed for direct commercial advantage that include this notice are permitted without fee. To copy otherwise or republish requires specific written permission.